

EXHIBIT 4

REDACTED

HIGHLY CONFIDENTIAL

**IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF VIRGINIA
ALEXANDRIA DIVISION**

United States of America, *et al.*,

Plaintiffs,

v

Google LLC,

Defendant.

Case No. 1:23-cv-00108

HON. LEONIE H. M. BRINKEMA

EXPERT REPORT OF
GABRIEL WEINTRAUB, PH.D.

DECEMBER 22, 2023

HIGHLY CONFIDENTIAL

II.A.2. Sell Side of the Ad Tech Stack

36. The end sellers of display advertising are publishers who originate internet content. One of the ways publishers monetize their work is by selling ad space on their websites.⁵¹ Like advertisers, publishers rely on intermediaries to procure programmatically the best ad that will generate the best financial or other results for the publisher that owns the website.
37. At a high level, publishers' goals include maximizing the yield (i.e., revenue) from their ad inventory. Publishers maximize their ad revenue by using the sell side ad tech tools to manage and optimize their yield. This includes analyzing the performance of different types of ads, e.g., guaranteed versus remnant ad impressions.⁵²
38. The main ad tech products that publishers utilize to sell their ad inventory programmatically⁵³ include PASs, ad exchanges, and SSPs, of which the latter two have, overtime, become interchangeable, as described further below.⁵⁴
39. Publishers use a PAS to serve ads on their webpages, accept ads of various types to fill inventory, and to measure the performance of the ads using metrics such as clicks and views. PASs also manage the sale of impressions across different demand sources such as ad networks and ad exchanges and also across direct and indirect channels.⁵⁵ In practice, a PAS accomplishes this by overseeing demand line items. Line items are the unit that PAS use when delivering an

⁵¹ at 3 (“The throttling mechanism paces the spending for a given advertising campaign, by calculating a participation probability for the campaign for each auctioned impression. This probability is calculated by the platform’s ad-server using real-time bidding information, such as the recent spending rate, remaining budget, remaining campaign time etc.”).

⁵² See Lee Report, Section II.A.1.

⁵³ “Get comprehensive yield management with Google Ad Manager,” Google Ad Manager, accessed December 13, 2023, <https://admanager.google.com/home/resources/feature-brief-yield-management/> (“Ad Manager helps you maximize your yield with a single, unified first price auction across all of your advertising demand. It compares the prices from your guaranteed campaigns with all non-guaranteed advertising sources – including real-time bidding partners, such as Authorized Buyers and Exchange Bidding partners – and prices from non-guaranteed line items, like those from your header bidding implementation, if you have one.”); see also, Lee Report, Section II.B.1.

⁵⁴ Ads sold “programmatically” are those that are sold using technology that automates and optimizes the ad buying process. See Lee Report, Section II.A.4.

⁵⁵ See, e.g., Lee Report, Section II.B.3.

⁵⁵ See Lee Report, Section II.B.1.

HIGHLY CONFIDENTIAL

71. For an auction platform to be viable, it typically needs to have thickness on both sides of the platform. A competitive and viable ad exchange, for example, should secure the participation of many publishers to provide abundant impressions to attract advertisers; at the same time, it should secure the participation of many advertisers to provide sufficient demand for impressions available on the platform to attract publishers.¹⁵⁸ The reason is that publishers and advertisers prefer an auction platform where they can find for each impression the best match possible, and thickness improves the “match quality” of transactions.¹⁵⁹ In open-web display ads, match quality can be measured by the advertiser’s valuation of the impression.¹⁶⁰ Advertiser spend or

¹⁵⁸ Zikun Ye, Dennis J. Zhang, Heng Zhang, Renyu Zhang, Xin Chen, and Zhiwei Xu, “Cold Start to Improve Market Thickness on Online Advertising Platforms: Data-Driven Algorithms and Field Experiments,” *Management Science* 69, no. 7 (2023): 3838–3860, at 3839 (“with higher market thickness, on one hand some user impressions that would otherwise be left unmatched can be matched with suitable ads and, on the other hand, the ads have more intensive competitions in the auctions on the platform.”); *see also*, Deposition of Benneaser John (Microsoft), September 8, 2023, 75:19–76:17 (“Q. Now, why does investing in the buy side have the effect of scaling the Xandr marketplace? A. The advertising ecosystem starts with advertisers, right, so advertisers would want to reach; and they’re the ones paying the money and paying the check, so you need to have the demand to fuel the ecosystem on the platform. Q. Right. So if you invest in the buy side, and fuel the demand, that has the effect of scaling the rest of the marketplace?...Is that what you’re saying? A. That is one way to scale the platform, because the demand is the key for scaling the marketplace; and providing the better tools and better reporting and measurement technologies would enable the advertisers to use our platform. That is the, that is one of the key investment pieces for investing in the DSP.”); *see also* David S. Evans and Richard Schmalensee, *Matchmakers: The New Economics of Multisided Platforms* (Boston: Harvard Business Review Press, 2016), 71 (“Multisided platforms face another huge hurdle that is related to the coordination problem. They don’t just need some customers from all sides to show up. They need enough to show up for it to matter. Without enough customers on all sides, a matchmaker is offering a lousy product.”).

¹⁵⁹ Yue Wu, Kaifu Zhang, and V. Padmanabhan, “Matchmaker Competition and Technology Provision,” *Journal of Marketing Research* 55 (2018): 396–413, at 397 (“[M]ost matchmaking markets are characterized by strong network effects. The probability of finding a good match is typically higher in a larger community.”); *see also*, Deposition of Benneaser John (Microsoft), September 8, 2023, 229:9–230:13 (“Q....How, if at all, does scale affect whether it’s easy or difficult for a company to enter the ad tech business? A. So when it comes to scale, let’s use the example of user. On the buy side, let’s say an advertiser wants to advertise their product across the age group of 40 and 50 or across a specific region. They won’t be able to reach through only one publisher or through only one farmer; so they are looking for multiple publishers where the user swings or goes or navigates multiple properties. So that’s the - that is the scale of reach advertisers need. And on the publisher side they don’t have the scale that they’ll be able to meet the needs of advertisers, so they’re looking for diversity of demand that they’ll be able to attract, and how will they participate into that market or network where they don’t have all the skills an advertiser is looking for.”).

¹⁶⁰ L. Elisa Celis, Gregory Lewis, Markus Mobius, and Hamid Nazerzadeh, “Buy-It-Now or Take-a-Chance: Price Discrimination through Randomized Auctions,” *Management Science* 60, no. 12 (2014): 2927–2948, at 2927 (“[Ad tech companies] also know more and more information about their users. This allows them to match advertisers to potential buyers with ever greater efficiency....For example, a sportswear firm advertising on the *New York Times* website may be willing to pay much more for an advertisement placed next to a sports article than one next to a movie review. It might pay an additional premium for a local consumer who lives in New York City and an even higher premium if the consumer is known to browse websites selling sportswear.”); *see also*, Deposition of Brian O’Kelley (AppNexus), September 29, 2023, 120:20–121:1 (“A. The whole idea of an auction is the highest price should win. And the more competition, the – you know, the more demand, the

HIGHLY CONFIDENTIAL

publisher's revenue can be a proxy for the valuation (even though these monetary quantities are also affected by other factors, such as competition, as I discussed above). For advertisers, this means buying impressions that connect them to end customers who value and engage with their ads, and for publishers, this means selling impressions to advertisers that have the highest willingness to pay for them.¹⁶¹

72. On the flip side, an auction platform that is denied access to scale in either advertisers or publishers will struggle to thrive even though it may offer an attractive technology. Economic principles of two-sided markets (as discussed above) imply that, because publishers hesitate to switch without access to sufficient demand, and advertisers do not want to try a service where there are not sufficient publishers to hit their ad spend goals, a new player will find it hard to achieve the level of scale and thickness needed on both the buy side and the sell side.¹⁶²

higher the price."); *see also*, Hana Choi, Carl F. Mela, Santiago R. Balseiro, and Adam Leary, "Online Display Advertising Markets: A Literature Review and Future Directions," *Information Systems Research* 31, no. 2 (2020): 556–575, at 557 ("[A]nd on the other side, publishers with consumers' impressions purvey ad inventory to advertisers with the highest valuations for those impressions. In between, intermediaries facilitate the match between advertisers and publishers by managing data and providing optimization tools and algorithms for serving ads.").

¹⁶¹ L. Elisa Celis, Gregory Lewis, Markus Mobius, and Hamid Nazerzadeh, "Buy-It-Now or Take-a-Chance: Price Discrimination through Randomized Auctions," *Management Science* 60, no. 12 (2014): 2927–2948, at 2930 ("These publishers earn revenue by selling advertising slots on these sites....Advertiser demand for each impression is determined by which user they are reaching, and what the user's current desires or intents are. For example, a Ferrari dealer might value high-income users located close to the dealership. A mortgage company might value people reading an article on 'how to refinance your mortgage' more than those reading an article on 'how to survive your midlife crisis,' whereas the dealership might prefer the reverse.").

¹⁶² David S. Evans and Richard Schmalensee, *Matchmakers: The New Economics of Multisided Platforms* (Boston: Harvard Business Review Press, 2016), 71 ("A multisided business, on the other hand, usually confronts the chicken-and-egg puzzle right away....There's a 'coordination problem.' Neither group will agree to use the service unless the other group does too."). Testimony from industry participants is consistent with this view. *See, e.g.*, Deposition of Benneaser John (Microsoft), September 8, 2023, 228:3–229:8 ("Q. So let me ask you about the US. In the US, how easy or difficult is it for new players to enter ad tech markets and compete effectively?...A. In the active market this is a broad - into the buy side and sell side and by different formats. As I said before, on the publisher side they're very hard to switch because of the demand access and the migration and scale to maintain. And on the buy side, advertisers want to go spend the dollars in the ad reach, and the users are more and more using services that offer, that's where the users go. And advertisers are not going to go somewhere else that they won't be able to see the return on ad value, on ad spend and to either build a business on a platform to get more advertising dollars or to build a publisher technology solution that's going to be, it's really, really hard. And that's why there are not many players that build ad servers, and tech companies and platforms are trying to come up with different ideas of where they can just like into the ecosystem and just build one-off solutions.").

HIGHLY CONFIDENTIAL

117. First, for fixed costs such as research and development investments, greater scale in terms of auction outcomes increases profitability. All else equal, higher number of impressions or higher revenue per impression will increase revenue without changing fixed costs. As an illustrative example, suppose that an ad tech company invested \$50 million in fixed cost. If the company won 10 percent of impressions that translates to \$100 million in net revenue (i.e., the revenue share to the company), then the company's profit would be \$50 million. As an illustrative example, if the company won 20 percent of impressions that translated to \$200 million, the company's profit would be \$150 million. If the company won only 1 percent of impressions, translating to \$10 million, then the company would have negative profits.
118. Second, for costs that vary with the volume of queries, such as server and processing costs discussed above, greater scale in terms of win rate (i.e., share of impressions won among queries) and revenue per impression increases profitability. To see this, note that profit can be written as revenue minus costs, where revenue is broken down into three parts as shown below:

$$\text{Profit} = \frac{\text{Revenue Per Impression} \times \text{Win Rate} \times \text{Queries}}{\text{Revenue}} - \frac{\text{Incremental Cost} \times \text{Queries}}{\text{Costs}}$$

The three parts of revenue are: (1) revenue per impression (i.e., $\frac{\text{Revenue}}{\text{Impressions}}$); (2) win rate, which is the number of impressions won among the number of queries that the ad tech product processed (i.e., $\frac{\text{Impressions}}{\text{Queries}}$); and (3) number of queries (*Queries*). Costs are calculated as the product of the incremental cost per query (*Incremental Cost*) and the number of queries. If the number of queries increases without a corresponding increase in either revenue per impression or the number of impressions won, then costs will increase but revenue will remain the same; thus, the company's profitability will decrease. However, for a given number of queries and revenue per impression, the ad tech company's profitability will increase as the win rate increases.

119. Importantly, these costs are incurred even if a platform fails to win the impression. Since ad exchanges, DSPs, and ad networks generate revenue when they win,³⁰⁰ platforms enjoying a

³⁰⁰ See, e.g., Deposition of Andrew Casale (Index Exchange), September 26, 2023, 174:6–14 (“Q. Now, does Index incur operating costs even if it fails to win an impression? A. Absolutely one hundred percent. Q. And when does Index generate revenue? A. Index only generates revenue if we clear a transaction or if an auction wins,

HIGHLY CONFIDENTIAL

likely worse off because another exchange would be able to provide a higher quality match than AdX or because another exchange may provide a similar match to AdX but extract lower fees from the transaction.

V.C.2. Google's Restriction of Real-time Bidding on AdX to DFP Prevented Rival Publisher Ad Servers from Attracting Publishers and Thereby Disadvantaged Rivals in Building a Viable PAS Business

161. Below, I highlight evidence from rival ad tech companies, particularly those that either currently have or used to have a PAS product, showing that Google's restriction of AdX's real-time bids to DFP—coupled with the restriction of Google Ads' demand to AdX—reduced the number of rival PAS platforms. Furthermore, evidence indicates that because rival PASs could not convince publishers to either add a second PAS or switch away from DFP and its “unique” demand source (Google Ads demand available through AdX), some rival PASs were unable to continue their operations.

162. Multiple industry participants emphasized that exclusive real-time access to Google Ads demand was a primary reason for publishers to use DFP. For example, John Gentry of OpenX observed that publishers had very limited access to AdX if they did not use Google's DFP as their publisher ad server and so it would be “economically unwise” for them to go with a different PAS that would “not have access to the largest pool of demand.”³⁷⁴ Arnaud Creput—CEO at Equativ³⁷⁵—pointed out that “a large part of Google demands, so from DSPs and Google Ads and DV360 is not accessible the same way by other SSPs. So if a publisher switched to a rival ad

attractive to sellers....This greatly weakens GDN's position in the market – why would an advertiser buy thru GDN when they can do the same retargeting on Criteo on same [sic] inventory *plus more inventory.*”)
(emphasis added).

³⁷⁴ Deposition of John Gentry (OpenX), October 26, 2023, 15:10–16:5 (“Q. How, if at all, did the fact that publishers had zero access to AdX if they did not use Google's DFP as their publisher ad server affect OpenX's ability to compete with DFP in the publisher ad server business?...[A.] It prevented our ability to compete in the publisher ad server business over time....Q. What do you mean by that? A. Fundamentally, the inability for OpenX to offer an ad server that could access the largest source of advertiser demand in the Internet, which was represented by AdX, imagine if a publisher went with the OpenX ad server, they were going to have a much lower amount of potential revenue that they could acquire and, as a result, much less monetization. So it was economically an unwise choice for them to choose to go with someone that did not have access to the largest pool of demand.”).

³⁷⁵ Deposition of Arnaud Creput (Equativ), September 5, 2023, 7:21–23 (“Q[.] So I think you are the CEO of Equativ, is that right? A[.] Right.”).

HIGHLY CONFIDENTIAL

server, he accepts actually to lose a significant part of its digital advertising revenues.”³⁷⁶ AdX’s demand is an important consideration in publishers’ decision of which PAS to use, and publishers use DFP over rival publisher ad servers because of the link between AdX and DFP; which publishers describe as “intertwined.”³⁷⁷ Publishers like News Corp confirmed that access to AdX demand was the primary reason they would not switch away from Google’s ad server.³⁷⁸

163. Several of Google’s competitors identified the connection between DFP and AdX demand as a key reason why publishers have been unwilling or unable to switch away from DFP to their platforms. For example, Brian O’Kelley of AppNexus relayed that publishers told AppNexus that they would not switch PAS platforms because of the risk of losing AdX demand.³⁷⁹ When asked why Xandr’s publisher ad server has had “no success for several years” in the U.S.,

³⁷⁶ Deposition of Arnaud Creput (Equativ), September 5, 2023, 17:25–19:2 (“A....[I]t was never about product pictures. It was never about level of service, which are considered by our clients as much better than the one of Google. But it was always for two reasons. Number one, switching cost. Google makes it very - (undecipherable) - actually to switch from their solution to other solutions. And number two is the fear of losing revenues since the fear not to have access to a Google Demand. So first to Google AdX. Google AdX is not available with the rival ad servers, except with a workaround, which is very imperfect, which is called AdX Mediation. It’s totally imperfect and it’s not working the same way with other SSPs how integrated with publishers. And second is the fact that a large part of Google demands, so from DSPs and Google Ads and DV360 is not accessible the same way by other SSPs. So if a publisher switched to a rival ad server, he accepts actually to lose a significant part of its digital advertising revenues.”).

³⁷⁷ Testimony of third-party market participants corroborate this point. *See, e.g.*, Deposition of Michael Shaughnessy (Kargo), August 9, 2023, 43:24–44:10 (“Q. So what role, if any, does AdX’s demand play in publisher’s decisions of which publisher ad server to choose?...A. It is a significant consideration for publishers. Q. And what’s your basis for saying that? A. I’m basing that on my experience as a publisher who is responsible with making these types of decisions.”); *see also, e.g.*, Deposition of Michael Shaughnessy (Kargo), August 9, 2023, 21:7–22:13 (“Q. How, if at all, does AdX’s scale, in terms of Open Auction display advertising transactions, give it competitive advantages over other SSPs?...A. Its scale and its ability to match with inventory, in my estimation, is higher than most. Q. And why do you say that?...A. Based on how DFP and AdX are intertwined. Q. What do you mean when you say AdX and DFP are ‘intertwined’? A. Publishers believe that they need to use DFP in order to leverage the demand of AdX. Q. And do you believe it’s true that publishers need to use DFP to leverage the demand of AdX? A. When I was a - ...When I was a publisher, that was my belief and how I operated. Q. And what was your basis for reaching that conclusion? A. Based on how AdX and pricing rules and, at one point, I believe AdX removed some of the tag functionality of how you would traffic within the ad server. So that is my recollection.”).

³⁷⁸ Deposition of David A. Minkin (News Corp) September 22, 2023, 190:15–20 (“Q....What were the reasons that News Corp decided not to switch to AppNexus and instead remain with GAM? A. It was primarily because of the unique programmatic demand that is only accessible via GAM.”).

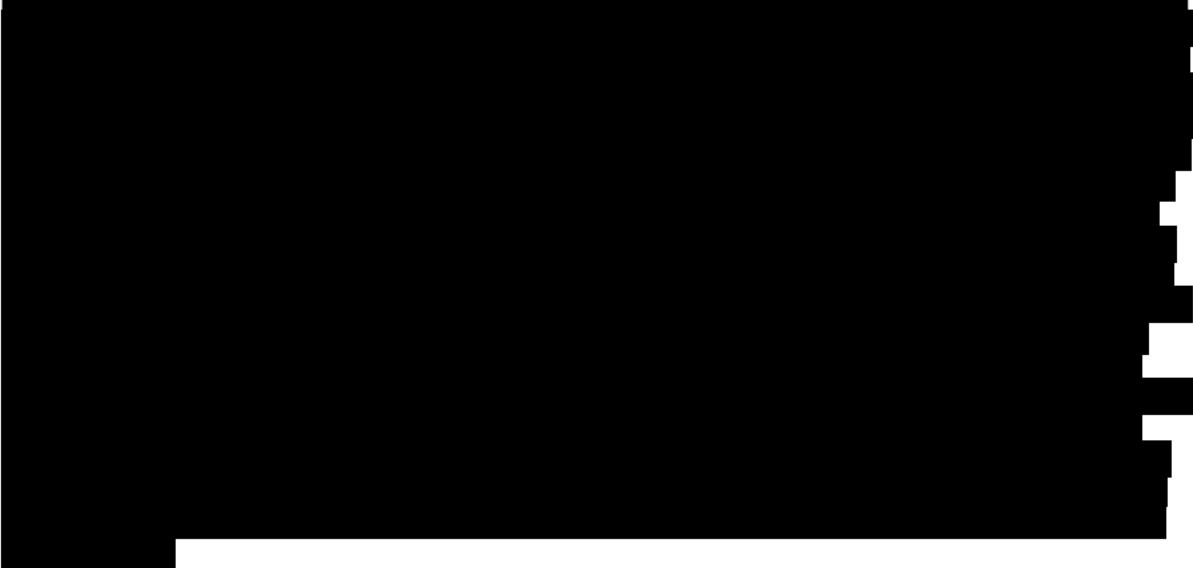
³⁷⁹ Deposition of Brian O’Kelley (AppNexus), September 29, 2023, 117:21–118:8 (“A. Because AppNexus built a competitive ad server, we would go into pitches with publishers, and the publisher would say that. They would say, you know, we are afraid of losing revenue from Google, if we switch. And we’ve talked to the Google team, and they’ve told us, you know, that that very well may happen. That because AdSense only buys through AdX, you know, leaving DFP would mean losing access to AdSense revenue.”).

HIGHLY CONFIDENTIAL

Benneaser John of Microsoft responded that “[n]umber one is the AdX demand” where “[publishers] are worried about losing the Google AdX demand.”³⁸⁰ Arnaud Creput of Equativ provided an example where, in 2018, Equativ endeavored to win clients from Google and was unsuccessful. Mr. Creput named the “fear not to have access to a Google demand,” as a reason for the failed attempts.³⁸¹

164. Switching costs also can be prohibitive for publishers to even consider switching PAS as Mr. Creput of Equativ confirmed as a key reason why publishers chose not to leave Google and that “it was never about the level of service.”³⁸² There are multiple reasons why the costs of switching a PAS are high. One reason is that the time, money, and effort associated with setting up the infrastructure for a publisher to work with a new ad server is a deterrent to switching or

³⁸⁰ Deposition of Benneaser John (Microsoft), September 8, 2023, 159:14–23 (“Q. Why did you say that Xandr Publisher ad server has been trying to penetrate the US market with no success for several years? A. Number one is the AdX demand. Number two is migration is very hard. Publishers had to put an effort to migrate. And the primary reason is the demand. They are worried about losing the Google AdX demand.”).

³⁸¹ A large rectangular area of the page has been completely blacked out, indicating redacted content.

³⁸² Deposition of Arnaud Creput (Equativ), September 5, 2023, 17:25–19:2 (“A....[I]t was never about product pictures. It was never about level of service, which are considered by our clients as much better than the one of Google. But it was always for two reasons. Number one, switching cost. Google makes it very - (undecipherable) - actually to switch from their solution to other solutions. And number two is the fear of losing revenues since the fear not to have access to a Google Demand. So first to Google AdX. Google AdX is not available with the rival ad servers, except with a workaround, which is very imperfect, which is called AdX Mediation. It’s totally imperfect and it’s not working the same way with other SSPs how integrated with publishers. And second is the fact that a large part of Google demands, so from DSPs and Google Ads and DV360 is not accessible the same way by other SSPs. So if a publisher switched to a rival ad server, he accepts actually to lose a significant part of its digital advertising revenues.”).

HIGHLY CONFIDENTIAL

adding a new PAS.³⁸³ Another reason is that the system of managing inventory and information associated with the inventory (that often would be captured in a “tag”) differs.³⁸⁴ Internal Google documents recognized that ad servers tend to be “sticky.”³⁸⁵ Moreover, publishers are unlikely to use simultaneously multiple ad servers for display advertising because, as described by James Glogovsky of The New York Times, it is “operationally challenging to rotate between two different ad servers for the personnel and teams responsible for trafficking and monitoring the campaign performance for our advertisers.”³⁸⁶

³⁸³ Deposition of Ryan Pauley (Vox), August 23, 2023, 95:16–96:6 (“Q. For a publisher, how easy or difficult is it to switch publisher ad servers?...A. It’s generally pretty difficult. It requires a lot of time and people. Q. And why is that? A. It’s a relatively complicated process to migrate from an existing ad server to implementing a new ad server, setting up the ad server, migrating any ongoing campaigns, the actual switching of the ad calls on your websites to a new ad server, things like that.”).

³⁸⁴ Deposition of Benneaser John (Microsoft), September 8, 2023, 179:20–180:24 (“Q. Why is migrating publisher ad servers difficult?...A. So if you’re a publisher, it starts with auto management system. So each publisher, they have their own auto management system, how they manage their inventory, which page and how the prices get sold and whatnot. And they have different channels of sales, direct sales, programmatic PG and all of those things. And then, when you look at the pages, there are tags in each of the pages, and each tag has a different settings on float rise and optimization and calling the measurements; so there are so many things that they need to manage and change at the page level. So there are three types of work involved. One is at the sales level on sales channels and whatnot. The second one is at the AdX level, how do they operate the inventory and how do they manage the placements and all of that at AdX level. The third one is anything at the tag level or the page level engineering team or ID teams need to change. So it’s like a complete change of the properties to migrate from one ad server to another ad server.”); *see also*, Deposition of Michael Shaughnessy (Kargo), August 9, 2023, 41:20–43:6 (“Q. And why do you believe that it’s not an area where you can compete?...A. There has been significant investment by both Google and publishers to create an ad server that works well, and many publishers use it for their direct sold campaigns. And even more important to them, as I understand, is being able to leverage the AdX demand within the Google ad server. Q. You said that it requires - there was investment by publishers; is that right? A. Yes. Q. What is that investment by publishers you were referring to? A. Over the years, publishing teams have had to build out their taxonomies, so the way that they classify the pages and how the ad units are set up. So that’s an infrastructure thing, which requires a developer as well as coordination with the ad operations team as well as their sales force, that they have a direct sales team. There also is the same infrastructure that is traditionally the foundation for all of the yield management when you’re working with other SSPs. There are reporting capabilities as well. So there are a lot of things that go into delivering an ad on a publisher’s page. Q. And if a publisher was to switch publisher ad servers, it would have to redo all of that investment; is that right? A. They would.”)

³⁸⁵ GOOG-TEX-00076049, at -049 (09/10/2012) (“Ad Servers are sticky, and hard to replace.”); *see also*, GOOG-TEX-00109645, at -647 (07/25/2019) (“We have long held the belief that ad servers are sticky and that by ‘owning the tag’ we would be best positioned to meet the needs of publishers, advertisers, and users more effectively. That theory has worked well for us.”); *see also*, GOOG-DOJ-01439665, at -669 (02/11/2009) (“Question: Is it sustainable to be ahead on the exchange if you don’t have the platform? David: My view is nothing really matters but the platform. Nothing has such high switching costs. If there’s a better network or exchange, you can just switch to it. Switching platforms is a nightmare. Takes an act of God to do it.”).

³⁸⁶ Deposition of James Glogovsky (The New York Times), August 25, 2023, 266:18–267:5 (“Q. Sure. How easy or difficult would it be for The New York Times to use more than one publisher ad server for display? A. It would be operationally challenging to rotate between two different ad servers for the personnel and teams responsible for trafficking and monitoring the campaign performance for our advertisers. It would also be

VI.A. Google's Implementation of Dynamic Allocation's First Look and Last Look Favored AdX to Win Over Rival Exchanges

VI.A.1. Google's Dynamic Allocation and How It Was Enabled by Google's Ownership of Ad Tech Products in Multiple Layers of the Ad Tech Stack

179. Dynamic Allocation is the process by which Google determined how remnant ad space is sold through the indirect sales channel.⁴⁰⁵ Specifically, this process uses an ad inventory allocation mechanism called the waterfall. In Google's waterfall, Google's PAS—DFP—calls ad exchanges in descending order based on their historical bids, often taking on the values of the exchanges' average historical CPMs.⁴⁰⁶ Notably, prior to Header Bidding and Open Bidding, DFP placed Google's own exchange, AdX, at the top of the waterfall over rival exchanges as its default setting, and I refer to this placement as "first look." With "first look," AdX had the opportunity to submit a bid for the inventory before other exchanges in the waterfall.⁴⁰⁷
180. Moreover, Google's "first look" placement allowed AdX to submit a real-time bid to compete for the ad space whereas other exchanges were ranked based on their historical average bids.⁴⁰⁸ Specifically, when DFP selected the best non-guaranteed line item to fill publisher inventory (and if the impression had not yet been served by a line item with a higher priority than AdX), it called a real-time auction for that inventory in AdX.⁴⁰⁹ If AdX's real-time price beat the price of

⁴⁰⁵ Google introduced "Enhanced Dynamic Allocation" in March 2014 where DFP filled some impressions through the indirect channel, rather than direct channel, when it determined that the direct contract could be met and a higher bid was available through the indirect channel. Today, Google refers to both Dynamic Allocation and Enhanced Dynamic Allocation as "Dynamic Allocation." GOOG-DOJ-05782415, at -431 (11/22/2019) ("Dynamic Allocation was improved in March 2014 (to create Enhanced Dynamic Allocation, now simply referred to as **Dynamic Allocation**)...In...2014, Enhanced Dynamic Allocation was introduced to enable publishers to allow remnant line items and Google's own line items – such as AdX line items and AdSense line items – to compete simultaneously with the guaranteed reservation line items with no impact to the delivery of reservation line items.").

⁴⁰⁶ For a more detailed discussion of Google's waterfall, *see* Ravi Report, Section III.A.

⁴⁰⁷ *See* Ravi Report, Section III.A.1.a.

⁴⁰⁸ *See* Ravi Report, Section III.A.1.a.

⁴⁰⁹ *See* Ravi Report, Appendix C.2.

HIGHLY CONFIDENTIAL

competing line items,⁴¹⁰ which include non-AdX exchanges' prices based on their historical bids, and also beat the reserve price, then AdX won the impression.⁴¹¹

181. There are two key implications to Google's "first look" placement regarding AdX's advantages over the other exchanges. First, prior to the arrival of Header Bidding and Open Bidding, AdX consistently had the opportunity to bid on and win an impression before other buyers would be called by DFP.⁴¹² Second, AdX offered a "real-time" price, because it ran a contemporaneous auction.⁴¹³ In comparison, non-AdX bidders "competed on average price" even if their advertisers had real-time bids that were higher than their historical averages and possibly even higher than the real-time AdX bids.⁴¹⁴ Therefore, until the advent of Header Bidding and Open Bidding, AdX's competitors only had a chance to bid for an impression if the static price of its line item within DFP, often set based on average historical prices, was sufficiently high.⁴¹⁵
182. AdX additionally had what is referred to as the "last look" advantage. As far back as the late 2000s, DFP let AdX use the highest price among other bidders (that may include exchanges, ad networks, and DSPs eligible to submit remnant line items) in the same auction as a floor in its contemporaneous auction.⁴¹⁶ Effectively, this "last look" feature gave AdX an informational edge over rival exchanges.⁴¹⁷ Even with Header Bidding, which attempted to counter Google's

⁴¹⁰ Line items are the unit of delivery in publisher ad servers. Line items define the price paid for an impression, the creative to be served, and other ad campaign elements such as campaign goals, duration, targeting, etc. See Ravi Report, Appendix C.1.

⁴¹¹ See Ravi Report, Section III.A.1.a.

⁴¹² See Ravi Report, Section III.A.1.a.

Even after the arrival of Header Bidding, AdX maintained an informational advantage over Header Bidding through its last look advantage, which I discuss in more detail in the following paragraph; moreover, AdX developed algorithms such as sell side dynamic revenue sharing that took Google's last look advantage to increase AdX's wins (see Section VI.B below).

Although rival exchanges participating in Open Bidding could compete with AdX with real-time bids, adoption of Open Bidding was limited in part because of a 5 percent fee that Google charged. See ¶169 above.

⁴¹³ See Ravi Report, Section III.A.1.a.

⁴¹⁴ GOOG-DOJ-14156104, at -104 (09/19/2016) ("[T]hird-party networks competed on average price vs. AdX per-impression bids"); see also, GOOG-DOJ-AT-01917966, at -983 (06/03/2020) ("When Dynamic Allocation is enabled, Publishers using GAM configure what is known as the 'value CPM' for their remnant line items. One common formula for setting the value CPM is to divide total revenue from the remnant line item over a given period by total impressions sent by the ad server to that line item to calculate the average CPM...").

⁴¹⁵ See Ravi Report, Section III.A.1.a.

⁴¹⁶ See Ravi Report, Section III.A.1.a.

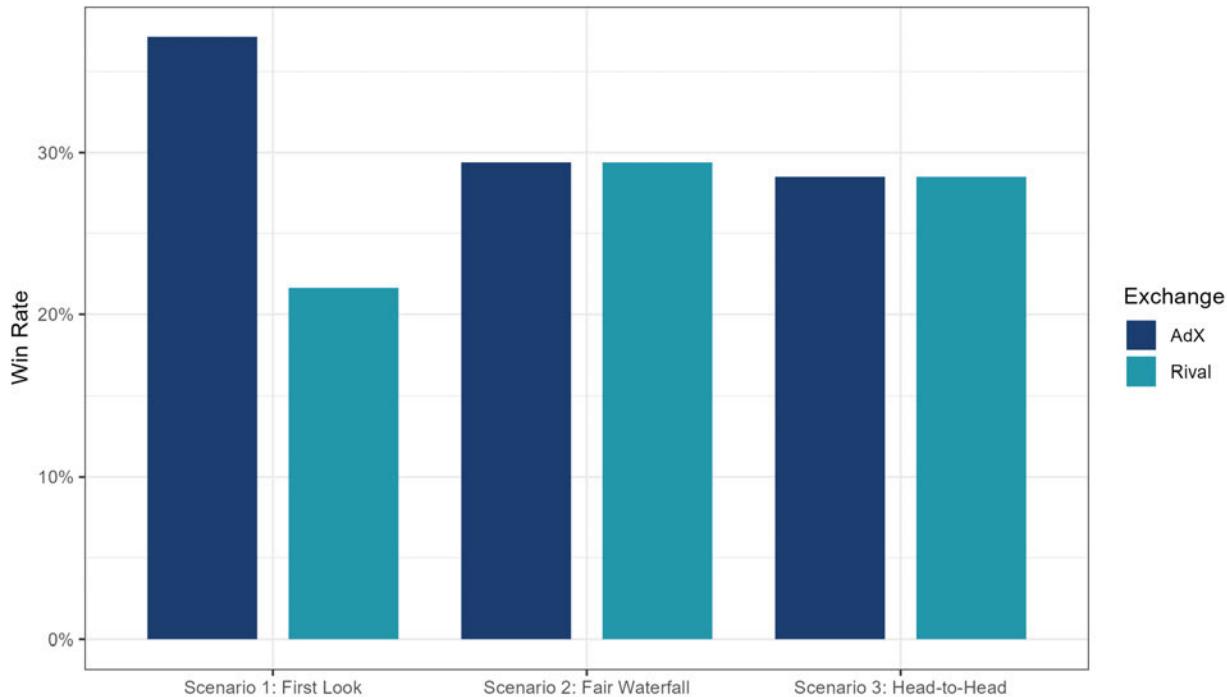
⁴¹⁷ See Ravi Report, Section III.A.1.a.

HIGHLY CONFIDENTIAL

the publisher's reserve price rather than beating the highest bid from the other exchange which could be higher than both the reserve price and AdX's highest bid.

190. My simulation results also show that AdX's advantage disappears in both Scenario 2 (the "fair" waterfall without "first look") and Scenario 3 (simultaneous head-to-head bidding) for a full range of the bid uncertainty. For the bid uncertainty of 0.65, AdX's probability of winning is equal that of the other exchange's probability of winning in both Scenario 2 and Scenario 3. This result is not surprising because the level of bids largely determines the winner of auctions in both Scenarios 2 and 3, and the only reason AdX wins more often in Scenario 1 is because it has the opportunity to bid before the other exchange as a result of "first look."

FIGURE 14: COMPARISON OF WIN RATES FOR ADX AND RIVAL EXCHANGE BY SCENARIO



Source: First Look Analysis Workpaper.

Notes: Spread = 0.65, Lambda = 1.

191. Regarding the impact on advertisers and publishers, I find that the overall advertiser surplus does not change very much across these three scenarios although AdX's advertiser captures a disproportionately larger portion of it under "first look" relative to the other two scenarios. Furthermore, I find in my model that "first look" yields suboptimal matches, that is the

HIGHLY CONFIDENTIAL

impression is misallocated (i.e., the impression is won by an exchange while the bidder on the other exchange values it more). For S=0.65, the probability of misallocation is equal to 7 percent. In this way, “first look” likely helped AdX’s advertisers win auctions even when advertisers bidding through rival exchanges had a higher value impression, which has the impact of hurting the publisher’s revenues as well as the overall efficiency of the marketplace. A marketplace that does not yield efficient allocations may discourage participation.

192. Regarding publishers’ revenues, as shown in Figure 15 below, the overall payment for the impression is higher for Scenario 3 (head-to-head competition) than it is for the other two waterfall scenarios. This is because while in the waterfall and “first look” scenarios the reserve price determines the payment, in the head-to-head scenario the second-highest bid determines the payment in cases where it exceeds the reserve price.
193. In addition, evidence from industry participants indicate that publishers make more money by using multiple exchanges that compete at the same time rather than sequentially, as is the setup in Google’s “waterfall.”⁴³⁹ Furthermore, publishers suffered under “first look” because it likely limited the choice of DFP publishers with regards to which exchange to sell their inventory. Google’s internal documents discuss the desire of publishers to diversify their activities across exchanges to increase their revenues. For example, in an internal Google presentation dated May 2016, Google employees notes that publishers “gave feedback that existing [AdX] solutions such as Enhanced Dynamic Allocation (EDA) or DFP First Look (DFL) does [sic] not address their need to increase yield with multiple exchanges.”⁴⁴⁰

⁴³⁹ Deposition of Tim Craycroft (Amazon and Google), August 15, 2023, 59:15–24 (“Publishers make more money when they use an ad exchange to supplement their direct-sold demands. They make more money when they use multiple ad exchanges. But the traditional way of using multiple ad exchanges is to give them each a try in sequence to serve an ad. Header bidding enables all those ad exchanges to compete at the same time, which enables better price discovery, which drives improved revenue for the publisher.”).

⁴⁴⁰ GOOG-DOJ-15277215, at -216 (05/05/2016) (“Publishers gave feedback that existing [AdX] solutions such as Enhanced Dynamic Allocation (EDA) or DFP First Look (DFL) does [sic] not address their need to increase yield with multiple exchanges”).